

D. White, R. Sharpe and N. Champagne

October 1, 2000





DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Work performed under the auspices of the U. S. Department of Energy by the University of California Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from the Office of Scientific and Technical Information P.O. Box 62, Oak Ridge, TN 37831 Prices available from (423) 576-8401 http://apollo.osti.gov/bridge/

Available to the public from the National Technical Information Service U.S. Department of Commerce 5285 Port Royal Rd., Springfield, VA 22161 http://www.ntis.gov/

OR

Lawrence Livermore National Laboratory Technical Information Department's Digital Library http://www.llnl.gov/tid/Library.html

Solution of Hybrid FEM-BEM Systems via Schur Complement Techniques

Daniel White, Robert Sharpe, Nathan Champagne October, 2000

1) Introduction

We are concerned with the numerical solution linear systems that arise from a hybridization of the Finite Element Method (FEM) and the Boundary Element Method (BEM). Our present focus is hybrid FEM-BEM discretization of the frequency-domain vector Helmholtz equation of electromagnetics, but similar hybrid techniques are used in electrostatics, acoustics, elasticity, etc. The hybrid FEM-BEM technique is used to solve "open" or "infinite" problems, where the FEM is used to discretize the interior of the problem and the BEM is used to simulate the effect of the infinite domain. This is illustrated generically in two dimensions in Figure 1 below. The FEM is applied to the interior V, the BEM is applied to the fictitious surface S, and the two methods are appropriately coupled to form a well-posed problem.

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

¹ A problem is well posed if 1) the solution exists, 2) the solution is unique, and 3) the solution depends continuously upon the source and/or boundary conditions. This is distinct from well-conditioned.

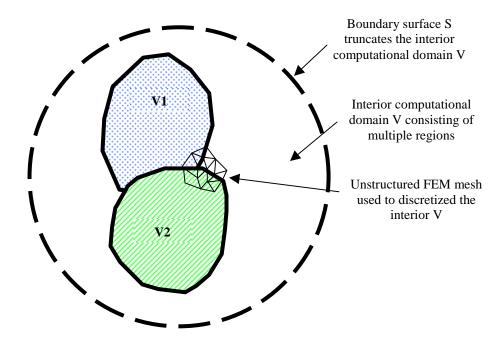


Figure 1. Computational volume discretized with a finite element mesh

As a concrete example, a standard FEM approximation for the frequency domain vector Helmholtz equation in an inhomogeneous volume V is given by

(1)
$$\sum_{V} \mu^{-1} \nabla \times W \cdot \nabla \times E^{\text{int}} dv - k^{2} \sum_{V} \mathcal{E}W \cdot E^{\text{int}} dv - ikZ \sum_{S} W \cdot \hat{n} \times H^{\text{ext}} ds = f^{\text{int}}$$

where E^{int} is the electric field in the interior, H^{ext} is the magnetic field on the surface, W is a finite element basis function (the test function), and f^{int} represents the application specific excitation. By itself equation (1) is of little use since H^{ext} is not known. A specific BEM known as the magnetic field integral equation can be used to determine H^{ext} ,

(2)
$$-\frac{1}{2} \sum_{s} Q \cdot \hat{n} \times H^{ext} ds - \sum_{s} Q \cdot \hat{n} \times \nabla \times G \times \hat{n}' \mathcal{G} \cdot H^{ext} ds' ds - ikY \sum_{s} Q \cdot \hat{n} \times G \times \hat{n}' \mathcal{G} \cdot E^{ext} ds' ds = f^{sur}$$

where E^{ext} is the electric field on the surface, Q is a finite element basis function, and G is the appropriate dyadic Green function. Different Green functions are used for different exterior regions, i.e. vacuum, layered media, random media, etc. Finally, equations (1) and (2) are coupled by equating E^{int} with E^{ext} on the surface,

(3)
$$\underbrace{\mathbf{l}}_{S} Q \cdot \hat{n} \times \mathbf{l} E^{\text{int}} - E^{\text{ext}} \mathbf{l} ds$$

Approximating the fields \boldsymbol{E}^{int} , \boldsymbol{E}^{ext} , and \boldsymbol{H}^{ext} in terms of basis function expansions

(4)
$$E^{\text{int}} = \sum_{i} E_{i}W_{i} \quad E^{\text{ext}} = \sum_{i} E_{i}Q_{i} \quad H^{\text{ext}} = \prod_{i} H_{i}Q_{i}$$

gives a system of linear equations Z x = y of dimension $N = N_{EV} + N_{ES} + N_{HS}$ where N_{EV} is the number of volume electric field unknowns, N_{ES} is the number of surface electric field unknowns, and N_{HS} is the number of surface magnetic field unknowns. The above formulation is described in greater detail in [1], and is an example of one of many different hybrid FEM-BEM formulations recently added to LLNL's EIGER code [2].

The system Z x = y is a hybrid sparse-dense linear system, the sparse part is due to the FEM and the dense part is due to the BEM. This matrix is complex valued, non-symmetric, indefinite, and ill-conditioned². The unknowns can be ordered such that the dense part is collected in the lower right corner as shown in Figure 2 below. Figure 2 is somewhat misleading because it is a relatively small problem; in general, as the mesh is refined the dense block becomes proportionally smaller. Naturally the precise ratio of the sizes of the dense and sparse blocks is problem specific, but for many problems of interest the ratio is determined by the surface-to-volume ratio of the mesh.

² The condition number of a matrix is $\chi \mathbb{I} = ||Z^1|| ||Z^{-1}||$ where $||\cdot||$ is your favorite consistent matrix norm. A matrix with large condition number is said to be ill-conditioned. The condition number is related to the spread in the spectrum of the matrix. Iterative methods typically take a very long time to converge for ill-conditioned systems, and typically converge quite rapidly for well-conditioned systems.

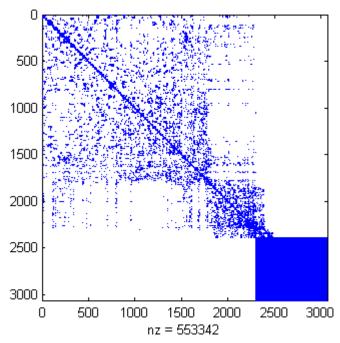


Figure 2. Sparsity pattern for an example hybrid FEM-BEM matrix. This particular problem had 2303 internal unknowns and 765 boundary unknowns.

At present the EIGER code supports two methods for solving linear systems: a direct LU decomposition [3] and the iterative BiCGStab method [4]. The direct LU method requires approximately $1/3 \text{ N}^3$ operations for a system of dimension N. This operation count is problem independent. The disadvantage of this direct method is that it does not take into account the structure of the matrix; it requires a large amount of memory N^2 and a large number of operations $1/3 N^3$ regardless of the relative sizes of N_{EV} , N_{ES} , N_{HS} .

The BiCGStab method is an example of a Krylov method, i.e. an iterative method that requires only the action of the matrix on a vector (or perhaps transpose operations as well). Krylov methods are ideal for sparse matrices as matrix vector multiplication requires only O(N) operations. If few iterations are required, this can be significantly more efficient than a direct method. Even if a relatively large number of iterations are required, a Krylov method is much more efficient than a direct solver in terms of memory usage.

We applied BiCGStab to two hybrid problems. Problem #1 is the so-called electromagnetic whistle, a hollow perfectly conducting sphere with a small

spherical aperture. This problem consisted of 2303 interior unknowns and 765 boundary unknowns. Problem #2 is a dielectric cone consisting of 1028 interior unknowns and 1344 boundary unknowns. Even though these problems are of very modest dimension, BiCGStab failed to converge for both problems³. The GMRES algorithm was converging, painfully slowly, and was terminated after 10000 iterations. The GMRES without restart and with full orthogonalization is guaranteed to converge (in exact arithmetic) in N iterations, however without restarting this method is no more efficient than a direct method. The QMR and CGS algorithms both converged for both test problems, although the convergence was much slower than desired. The convergence rates are shown in Figures 3 and 4.

There are several methods for improving the convergence of Krylov methods. One approach is preconditioning the linear system. Instead of solving the system Z x = y for x, the modified system L Z x = L y is solved for x, where the matrix L Z is better conditioned than the matrix Z alone. This is one of many possible preconditioning strategies. The difficulty is constructing a preconditioner that is both easy to compute and significantly improves the conditioning of the linear system. We tried several "black box" linear algebraic preconditioners such as Sparse Approximate Inverse and Incomplete LU [4], however these had little effect upon the convergence. We suspect that a special-purpose preconditioner that is based on the unique characteristics of out problem will be required.

Another approach to improving convergence is to re-formulate the problem. For example, different variational forms or different finite element basis functions often lead to matrices with improved conditioning. Due to time constraints this approach was not investigated.

³ While the convergence of BiCGStab is known to be somewhat erratic, for both test problems the error grew monotonically, which is unexpected. There could be a bug in the code. However, the BiCGStab code was tested on other complex valued, non-symmetric, indefinite matrices and it always converged. Thus more investigation is required.

2) Schur Complement

One possible approach for improving the convergence of Krylov methods when applied to hybrid FEM-BEM matrices is to use the Schur Complement. This approach is attractive because it takes advantage of the sparse-dense structure of the matrix, and it does not require a radical reformulation of problem. We begin by writing the system of equations in block form

$$\begin{pmatrix}
A & B \\
C & D
\end{pmatrix}
\begin{pmatrix}
Y_1 \\
Y_2
\end{pmatrix} = \begin{pmatrix}
Y_1 \\
Y_2
\end{pmatrix}$$

where A, B, and C are sparse blocks and D is a dense block. The subvector x_1 represents the internal unknowns and the subvector x_2 represents the boundary unknowns. We can express x_2 in terms of y_1 and y_2 as

(6)
$$x_2 = D^{-1} \mathbf{y}_2 - C y_1 \mathbf{y}_2$$

and we can then write for an equation as x_1

In equation (7) the matrix $S = (A - BD^{-1}C)$ is the Schur Complement. The idea is that the matrix D is small enough to be "inverted" directly and the sparse system $Sx_1 = b$ can be solved via a Krylov method. Using surface-to-volume ratio arguments, the Schur Complement approach has a memory savings of $N^{2/3}$ compared to using a direct solver on the whole system Z = y. As an example, for a computer with 10^9 words of memory, the direct LU approach restricts us to problems of order $N\sim30,000$ whereas with the Schur Complement approach we can solve systems of order $N\sim5,000,000$.

Of course it is necessary to investigate the conditioning of the Schur Complement. If it is as ill-conditioned as the original Z matrix we have not gained anything.

We begin by examining the spectrum of the complete Z matrix for problems #1 and #2. The spectra are shown in Figures 5 and 6 below. The main point is that the eigenvalues are not clustered. It is interesting to note that the ratio

 $\max(abs(\lambda))/\min(abs(\lambda))$, which is an often (mis)used estimate of the condition number, is 7.64e3 and 7.0e4 for Problem #1 and Problem #2, respectively. According to this measure these problems are not ill-conditioned. But it is obvious that the eigenvalues are far from being clustered and this has a significant effect upon the convergence of iterative methods.

We next examine the eigenvalue spectrum of the Schur Complement for Problem #1 and Problem #2. These are shown, along with the original spectrum, in Figures 7 and 8 below. The original spectrum is shown in blue, the Schur Complement spectrum in red. Note that the eigenvalues are clustered along the real axis. In fact, the spectrum of the Schur Complement is comparable to that of the matrix A itself.

3) Results, Conclusions, and Issues

The following Krylov methods were applied to the Schur Complement systems: BiCGStab, CGS, QMR, and GMRES. For both Problem #1 and Problem #2, BiCGStab failed to converge. However the other methods converged nicely. The results are shown in Figures 9 and 10 below.

For both problems the GMRES algorithm was ran with a restart interval of 100. In both cases the convergence of GMRES tapered off, which is interesting because it is often considered "the best" method. It is a very robust method, as indicated by the steady, but slow, convergence. The CGS method on the other hand was quite erratic, but it ultimately gave us the answer in the shortest amount of time. The QMR algorithm was in the middle.

It is difficult to make conclusions based on this limited set of experiments. While the Schur Complement approach works, it is far from a scalable algorithm. It is important to note that the Schur Complement approach can be used in conjunction with a fast multipole algorithm for the BEM [1]. Use of a fast multipole algorithm for the evaluation of D⁻¹ will save significant amount of memory and reduce the cost of each iteration. However the number of iterations is depended upon the conditioning of A, the discretized Helmholtz operator. Hence more effort is required in the area of preconditioning this operator.

To summarize:

- Direct LU decomposition works, but the $O(N^2)$ memory usage and the $O(N^3)$ operation count place a sever constraint on the size of the problem.
- Standard Krylov methods will require much less memory than direct LU decomposition. However, standard Krylov methods applied to two example hybrid FEM-BEM matrices were painfully slow. A deeper understanding of BEM and hybrid FEM-BEM is required to develop preconditioners for accelerating Krylov methods.
- The Schur Complement approach as described in this technical note is a hybrid approach well suited for hybrid sparse-dense matrices. The Schur Complement approach takes advantage of the block structure of the matrix, resulting in a significantly memory savings compared with the direct LU approach.

- The Schur Complement approach significantly reduced the number of Krylov iterations required for convergence. The eigenvalue spectrum of the Schur Complement is comparable to that of the A matrix itself. However, the approach is at best an $O(N^2)$ method. To reduce the iteration count further a preconditioner for the A matrix is required.
- The Schur Complement approach does not preclude the use of fast multipole methods for the BEM. A fast multipole method will further reduce memory usage and reduce the cost of each iteration.

4) The Software

In order to perform the computational experiments described above it was necessary to actually write software. Rather than write an application-specific "one-time-use-only" Schur Complement solver we decided to develop an application-independent solver framework that could meet the needs of EIGER and other unstructured-grid FEM-type codes for several years. The solver framework was to have:

- Support for both real-valued and complex-valued linear systems,
- Support for a variety of matrix storage formats (including hybrid formats),
- A variety of different Krylov methods,
- A variety of preconditioners,
- A dense direct LU decomposition,
- A sparse direct LU decomposition,
- Distributed memory parallelism (i.e. MPI),
- Extensibility the ability to easily extend the framework to include new types of matrices, new solvers, new preconditioners, etc.

While not complete, much progress was made towards the development of a solver framework. We did not start from scratch; we began with the ISIS++ library developed by Sandia [5]. In Figure 11 we illustrate the class hierarchy of the solver framework, annotating the initial ISIS++ classes, the classes added this year, and possible future classes. Note that Figure 11 presents a "high-level" view of the framework, there are several "low-level" classes that are not illustrated.

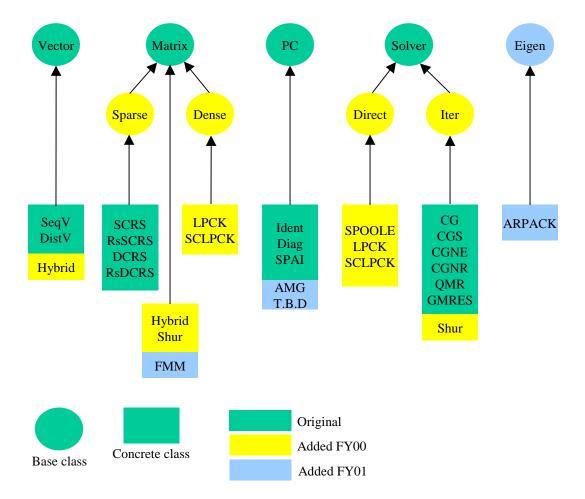


Figure 11. The class hierarchy of the solver framework. The basic abstractions are Vectors, Matrices, Preconditioners, and Solvers. Future abstractions will include eigenvalue solvers and non-linear equation solvers.

- 5) References
- [1] J. Volakis, A. Chatterjee, L. Kempel, Finite Element Method for Electromagnetics, IEEE Press, 1998.
- [2] R. Sharpe, J. Grant, N. Champagne, et. al., "EIGER: electromagnetics interactions generalized," IEEE APS Symposium, Quebec, Canada, July 13-17, 1997. UCRL-JC-130496.
- [3] G. Golub, C. Van Loan, <u>Matrix Computations</u>, Johns Hopkins University Press, 1989.
- [4] Y. Saad, Iterative Methods for Sparse Linear Systems, PWS Publishing, 1996.
- [5] ISIS++ Reference Guide, Sandia Report SAND99-8231, April 1999.

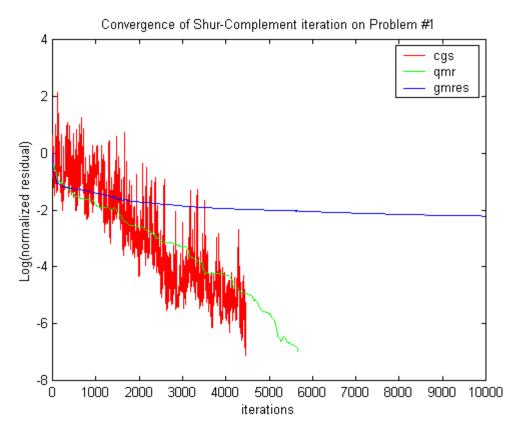


Figure 3. Convergence of various Krylov methods on test problems #1. The BiCGStab algorithm failed to converge. Diagonal scaling was employed. The problem dimension was 3068.

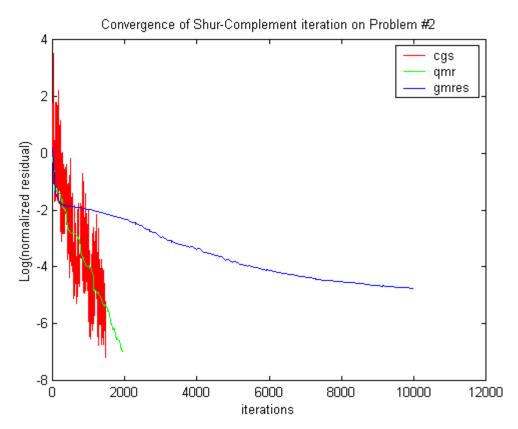


Figure 4. Convergence of various Krylov methods on test problems #2. The BiCGStab algorithm failed to converge. Diagonal scaling was employed. The problem dimension was 2372.

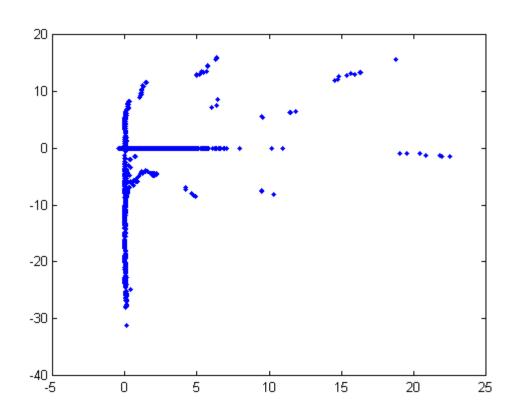


Figure 5. Eigenvalue spectrum of the Z matrix for Problem #1.

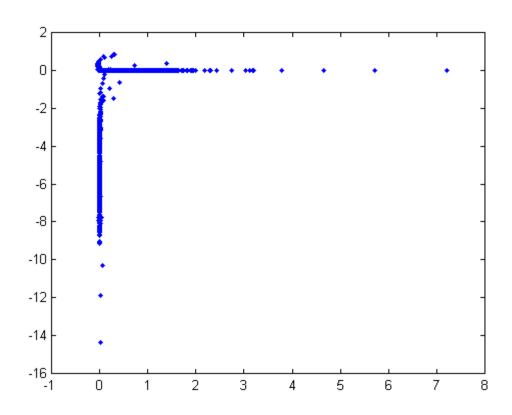


Figure 6. Eigenvalue Spectrum of the Z matrix for Problem #2.

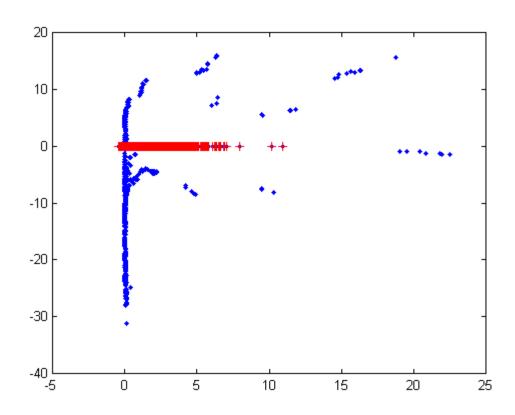


Figure 7. The spectrum of the Schur Complement vs. the spectrum of the original Z matrix for Problem #1. The spectrum of the Schur Complement resembles that of the FEM matrix A alone.

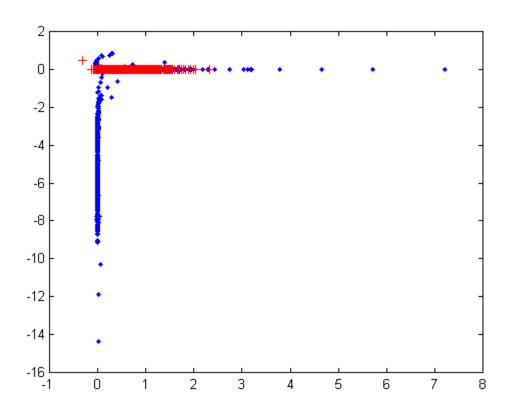


Figure 8. The spectrum of the Schur Complement vs. the spectrum of the original Z matrix for Problem #2. The spectrum of the Schur Complement resembles that of the FEM matrix A alone.

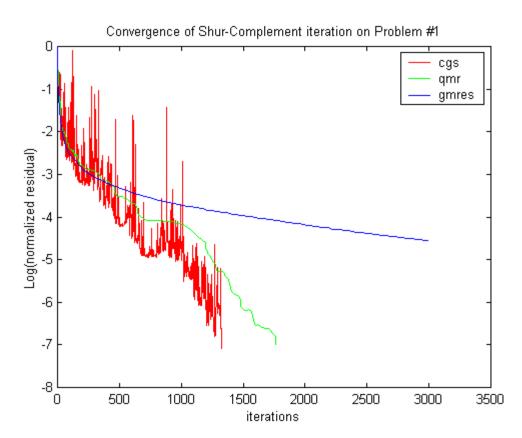


Figure 9. Converge of Schur Complement approach for Problem #1

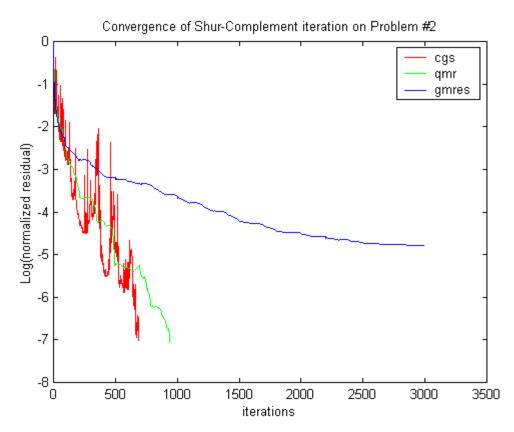


Figure 10. Convergence of Schur Complement approach for Problem #2.